

Simultaneous computation of dynamical and equilibrium information using a weighted ensemble of trajectories

Steven Lettieri¹, Matthew C. Zwier², Carsen A. Stringer³, Ernesto Suárez¹,
Lillian T. Chong², and Daniel M. Zuckerman^{1*}

1. Department of Computational and Systems Biology, University of Pittsburgh
2. Department of Chemistry, University of Pittsburgh
3. College of Arts and Sciences, University of Pittsburgh

* Corresponding author: ddmmzz@pitt.edu

October 12, 2012

Abstract

Equilibrium can be represented as an ensemble of uncoupled systems undergoing dynamics in which detailed balance is maintained. We describe a simulation protocol based on the “weighted ensemble” (WE) approach [Huber and Kim, *Biophys. J.*, 1996], in which weak coupling between parallel simulations can capture both dynamical and equilibrium properties without statistical bias. The approach is demonstrated on molecular systems. Of particular practical importance, analysis of the weighted ensemble does not require states to be chosen in advance, and permits the calculation of rates between arbitrary states chosen after the simulation is run. We describe a history-dependent non-Markovian formulation to correct for a potentially significant bias in kinetic properties measured from equilibrium WE simulations.

1 Introduction

Although it is textbook knowledge that the functions of biomacromolecules are strongly coupled to their conformational motions and fluctuations [1], computer simulation of such motions has been a challenge for decades [2]. Typically, distinct algorithms are employed to estimate equilibrium quantities (e.g., [3, 4]) and dynamical properties (e.g., [5–9]). In principle, a single long dynamics trajectory would be sufficient to determine both equilibrium and dynamical properties [10], but

such simulations remain impractical for most systems of interest. More technical approaches that can yield both equilibrium and dynamical simulation, sometimes under minor assumptions, have drawn increasing attention [11–15].

Here, we demonstrate a statistically exact simulation approach which directly embodies equilibrium (Fig. 1) via multiple dynamical replicas in detailed balance, but which also yields non-equilibrium dynamical information. The method generalizes the “weighted ensemble” (WE) strategy [5]. WE simulation has previously been used to characterize non-equilibrium dynamical pathways and rates in model molecular systems (e.g., [16–18]), and for equilibrium simulation [14, 19]. The rigorous path-sampling formalism underlying WE simulation has been described [20]. Very recent work pursues WE strategies similar to that proposed below [14, 15] without accounting for history dependence in the same fashion.

2 Theoretical formulation

WE simulation uses multiple simultaneous trajectories that are occasionally coupled by replication or combination events to gain efficiency over simple brute-force simulation [5]. The coupling events typically are governed by a static partition of configuration space into “bins” (Fig. 1c), although dynamical/adaptive bins may be used [20]. In the case of static bins, when one or more trajectories enters an unoccupied bin, those trajectories are replicated to conform to a (typically) preset value, M . Similarly, if more than M trajectories are found to occupy a bin, trajectories are combined statistically in a pairwise fashion until M remain. As described below, each trajectory is assigned a weight (i.e., probability) so that dynamics remain statistically unbiased [20]. Weights are reassigned at occasional intervals according to the condition being simulated (e.g., equilibrium or a nonequilibrium steady state) [19].

The goal of the present work is to reassign weights in a way that achieves equilibrium as rapidly as possible (after which rates between arbitrary states can be computed, as described below). Our procedure is based on the lack of probability flow among bins during equilibrium: $p_i^{\text{eq}} k_{ij}^{\text{eq}} = p_j^{\text{eq}} k_{ji}^{\text{eq}}$, where p_i^{eq} is the equilibrium probability (sum of trajectory weights) in bin i and k_{ij}^{eq} is the conditional transition probability per unit time (i.e., rate) from bin i to j observed in equilibrium. Early in a simulation, however, the lack of an equilibrium distribution implies that measured rates among the finite bins will differ from their equilibrium values. Nevertheless, we can assign non-equilibrium bin probabilities \hat{p} based on the transient, non-equilibrium rates \hat{k} according to the detailed balance condition

$$\hat{p}_i / \hat{p}_j \leftarrow \hat{k}_{ji} / \hat{k}_{ij} . \quad (1)$$

Alternatively, steady-state equations, which hold in equilibrium, can be solved for $\{\hat{p}_i\}$ values:

$$\sum_j \hat{p}_i \hat{k}_{ij} = \sum_j \hat{p}_j \hat{k}_{ji} . \quad (2)$$

Adjusting trajectory weights toward their equilibrium values via Eq. (1) or Eq. (2) is necessary for




Figure 1: Equilibrium in different representations. (a) Ensemble of trajectories with arrow tips indicating instantaneous configuration and tails showing recent history in the space of two schematic coordinates q_1 and q_2 . States A and B, shown in grey, are two arbitrary regions of phase space. (b) Dissection into two subsets based on whether trajectory was most recently in state A (black solid arrows, the “ α ” steady state) or state B (red dashed, “ β ”). (c) Statistically equivalent ensemble of weighted trajectories, with arrow thickness suggesting weight. Configuration space has been divided into cells (“bins”) which each contain an equal number of trajectories.

the rates \hat{k} to converge toward k^{eq} . Note that Eq. (2) tends to have a unique solution, unlike Eq. (1) which is overdetermined when rates \hat{k} are imperfectly estimated. As described below, however, the use of Eq. (1) or Eq. (2) is generally not sufficient to avoid bias in estimating *non-equilibrium* dynamical properties.

Dynamical quantities can also be calculated from an equilibrium WE simulation, with additional analysis. Once equilibrium *among trajectories* is established in a WE simulation, in the strict sense described below, the rates between arbitrary states A and B can be measured. Any region of phase space can be used to define a state, e.g., all configurations within a threshold “distance” from some reference configuration. Rates can be computed on-the-fly during a WE simulation if states are defined in advance, or in a post-analysis of saved trajectories. Several steps are required. First, as shown in (Fig. 1b), the equilibrium set of trajectories is decomposed into two steady states: the α steady state consisting of trajectories more recently in A than B, and the β steady state with those most recently in B [9,21]; these were denoted “AB” and “BA” steady states, respectively, in Ref. [21]. Trajectories can be “labeled” according to the last state visited — i.e., classified as α or β — during a WE simulation or in a post-analysis. See also Fig. 3.

Once a true equilibrium among trajectories is decomposed into the two steady states, a rate can be directly computed from the probability arriving to the final state [9, 19, 22] via

$$\frac{1}{\text{MFPT}(A \rightarrow B)} = \frac{\text{Flux}(A \rightarrow B|\alpha)}{p(\alpha)}, \quad (3)$$

where MFPT is the mean-first-passage time, $\text{Flux}(A \rightarrow B|\alpha)$ is the probability per unit time arriving to state B in the α steady state and $p(\alpha)$ is the total probability in the α steady state. By construction

$$\left[\begin{array}{ccc} k_{11} & k_{12} & k_{13} \\ k_{21} & k_{22} & k_{23} \\ k_{31} & k_{32} & k_{33} \end{array} \right] \quad \left[\begin{array}{cc|cc|cc} k_{11}^{\alpha\alpha} & 0 & k_{12}^{\alpha\alpha} & 0 & 0 & k_{13}^{\alpha\beta} \\ 0 & 0 & 0 & 0 & 0 & 0 \\ \hline k_{21}^{\alpha\alpha} & 0 & k_{22}^{\alpha\alpha} & 0 & 0 & k_{23}^{\alpha\beta} \\ k_{21}^{\beta\alpha} & 0 & 0 & k_{22}^{\beta\beta} & 0 & k_{23}^{\beta\beta} \\ \hline 0 & 0 & 0 & 0 & 0 & 0 \\ k_{31}^{\beta\alpha} & 0 & 0 & k_{32}^{\beta\beta} & 0 & k_{33}^{\beta\beta} \end{array} \right]$$

Figure 2: Constructing a labeled rate matrix for unbiased calculations. For purposes of illustration, here state A consists solely of bin 1 and state B solely of bin 2. Left: A traditional rate matrix with history-blind elements. The rate k_{ij} gives the conditional probability for transitioning from bin i to bin j in a fixed time increment, regardless of previous history. Right: The labeled rate matrix accounting for history. The element $k_{ij}^{\mu\nu}$ is the conditional probability for the i to j transition for trajectories initially in the μ sub-ensemble which transition to the ν sub-ensemble, where μ and ν are either α or β . The labeled rate matrix correctly assigns the α and β sub-populations of each bin, whereas the traditional matrix may not.

$p(\alpha) + p(\beta) = 1$. Normalizing by p effectively excludes the reverse steady state and the rate calculation only “sees” a uni-directional steady state as in Ref. [19]. An expression analogous to Eq. (3) applies for the $B \rightarrow A$ MFPT. Also note that the effective first order rate constant, defined by $\text{Flux}(A \rightarrow B|\alpha)/p_A^{\text{eq}}$, can be determined from equilibrium WE simulation because p_A^{eq} can be directly computed. Eq. 3 is the basis for the “direct” MFPT estimates reported below.

2.1 Correcting for bias in rate estimation

It is critical to realize that Eq. (3) assumes more than the configurational equilibrium obtained by application of Eq. (1) or Eq. (2). In fact, the simple application of Eq. (3) can lead to significant bias in rate estimates in equilibrium implementations of WE, as our data show. The error arises because the correct configurational distribution does not guarantee the correct distribution of α vs. β trajectories as required by Eq. (3). It is not clear whether this bias would be apparent in configurationally symmetric systems, but it is clear in the molecular examples studied below.

It is possible to construct a non-Markovian description that rigorously accounts for the α/β labeling of trajectories and yields unbiased rate estimates, even based on bins exhibiting highly non-Markovian behavior. The unbiased formalism is built from the observation that, as embodied in Eq. (3), the non-equilibrium (and generally non-Markovian) information required for rate estima-

tion is solely the α/β distribution. That is, the required trajectory history is whether a particular trajectory originated in A (an α trajectory) or B (a β trajectory) — and it is this property that can be biased by reweighting based on Eq. (1) or Eq. (2). The α (or β) designation is sufficient because then the α or β steady state is recovered precisely, leading in turn to the exact rates [9, 19].

To eliminate bias by strictly and separately conforming to the α and β steady states, we decompose the equilibrium population into α and β components for each bin i :

$$p_i^{\text{eq}} = p_i^\alpha + p_i^\beta, \quad (4)$$

which implies $p(\alpha) = \sum_i p_i^\alpha$ and $p(\beta) = \sum_i p_i^\beta$. We called this a “labeled” analysis. Thus, with N bins, a set of $2N$ probabilities is required rather than N . Similarly, a $2N \times 2N$ rate matrix is required: $k_{ij}^{\mu\nu}$, where μ and ν can be either the α or β subsets of trajectories. See Fig. 2.

Each of the previously considered k_{ij} rate elements is thus decomposed into four history-dependent elements which account for whether the particular trajectory was last in state A or B and whether the trajectory transitions between the α and β subsets. See Fig. 2. The analysis assumes states consist strictly of one or more bins, but this is always possible in a post-analysis without loss of generality. Note that more than half the $k_{ij}^{\mu\nu}$ elements are zero. For example, consider a bin in the ‘intermediate’ region (neither A nor B), such as bin 2 in Fig. 2: in this region an α trajectory cannot change into a β trajectory and *vice versa*; hence rates for these processes are zero. Similarly, an α trajectory in the intermediate region which enters a bin in B must turn into a β trajectory, so the rate will always be zero to the α components of bins in B.

The non-Markovian results below use the formalism stemming from Eq. (4). Rates $k_{ij}^{\mu\nu}$ are measured in a post-analysis and are analyzed in a Markov-like manner, despite storing potentially long-time history dependence via the α and β labels. Importantly, WE permits the history-dependent rates $k_{ij}^{\mu\nu}$ to be measured in a post-analysis of a simulation which was reweighted without regard to α/β histories, or even a simulation which was not reweighted at all. The results below are based on simulations reweighted via Eq. (2) but post-analyzed using the history dependence just described.

There are several advantages to a formulation using α/β labeled rates $k_{ij}^{\mu\nu}$. (i) Most importantly, bias is removed from rate estimates. (ii) Solution of both the α and β steady states is performed simultaneously via a standard Markov-state analysis of the $k_{ij}^{\mu\nu}$ rate matrix. By contrast, if the α and β steady states are independently solved within a Markov formalism, there can be substantial ambiguity in how to assign feedback from the target to initial state when the initial state consists of more than one bin. (iii) The labeled formulation guarantees, by construction, the flux balance intrinsic to equilibrium, namely, $\text{Flux}(A \rightarrow B|\alpha) = \text{Flux}(B \rightarrow A|\beta)$. (iv) The analysis can be performed using arbitrary bins (and states defined as sets of these bins). It is not necessary to employ the bins originally used to run the WE simulation because a post-analysis can calculate rates among any regions of configuration space.

In general, a post-analysis evaluation of kinetic properties is particularly powerful in permitting

the evaluation of different state definitions. Indeed, for complex systems, it is a major challenge to generate reasonable state definitions. In our approach, one can use kinetic properties determined in post-analysis of varying state definitions to optimally assign state definitions.

3 Model systems and simulation details

Equilibrium weighted ensemble simulations were performed on three systems of varying degrees of complexity: a pair of methane molecules, the small peptide alanine dipeptide (AD) and chignolin, a 10 residue mini-protein. All simulations were performed at 300K. For methane molecules, GROMACS united-atom interactions were used with TIP3P explicit solvent. For AD and chignolin, the all-atom parm99-SB force-field with implicit GBSA solvent was simulated in AMBER. The default GBSA parameters were used in AMBER version 11. The molecular dynamics time step used for all systems was $\Delta t = 2$ fs.

The reweighting procedure requires reliable estimates of the rate matrix in order to obtain unique and well-defined steady-state solutions. To achieve this, average rate matrices were used during the reweighting process. In this work, rate matrix averaging extended over the second half of all available simulation data at the time of the reweighting, although other schemes are possible. For all simulations, reweighting was performed after each iteration of the WE algorithm. An iteration is defined to be the simultaneous propagation of all trajectories in the ensemble for some amount of time, τ . In these studies, a value of $\tau = 250\Delta t$ is used.

For each system, progress coordinates were selected and “binned” in a 1D or 2D space. For methane, we used the distance r between the two methane molecules following [18]. The coordinate $r \in [0, 16]$ Å was partitioned with a bin spacing of 1 Å. For AD, the two dihedral angles $\phi \in [-180, 180]$ and $\psi \in [-180, 180]$ form the progress coordinate space. We used a 12 x 12 partition of the 2D progress coordinate space so that bins are 30° x 30°. For chignolin, two important folding-related distances were binned according to the recommendations in [23]. Specifically, the distances $X \equiv r(ASP3O - GLY7N) \in [0, 12]$ Å and $Y \equiv r(GLU5O - THR8N) \in [0, 12]$ Å were used with a bin spacing of 0.5 Å in each dimension.

To highlight the ability of equilibrium WE to compute both equilibrium and non-equilibrium properties, different observables are emphasized for each of the systems. For all systems, the potential of mean force (PMF) is calculated via $-\log(P(\mathbf{q}))$ where P is the probability distribution function and \mathbf{q} is the reaction coordinate vector (in units of $k_B T = 1$). Note that since trajectories have different weights, summing the weights of all trajectories in a given bin yields the probability of that bin. Repeating this for all bins yields $P(\mathbf{q})$. Since the PMFs are calculated in post-analysis, they can be measured at a finer resolution than the bin spacing used during the WE simulation. From $P(\mathbf{q})$, one can also infer the populations of arbitrary states since states are defined as sets of bins. Importantly, the PMF was used to aid in the selection of states for the MFPT analysis. The mean first passage times between arbitrary states can be measured either directly via Eq. 3 or via the

labeled rate matrix formulation.

For “labeled” non-Markovian rate estimation, the following procedure was used:

1. Estimate the labeled rate matrix $k_{ij}^{\mu\nu}$ using transitions occurring during WE simulation.
2. Determine the steady state solution. One way to accomplish this is to exponentiate the labeled rate matrix and multiply it by the vector $\mathbf{p}_0 = (1, 1, \dots, 1)/2N$ until it converges to the steady state solution \mathbf{p}_{SS} . Not all matrices will remain stable under exponentiation (i.e., their elements will approach zero). In such cases, the steady state solution is not well defined and measurements of the passage time are skipped. Typically, this only occurs in the beginning of simulations since new bins are being occupied for the first time and there are no estimates of rates out of said bins.
3. Use the steady state solution \mathbf{p}_{SS} along with the labeled rate matrix elements to calculate the α flux entering state B and the β flux entering A.
4. Finally, the MFPT values are given by Eq. 3 with fluxes computed according to

$$\text{Flux}(A \rightarrow B|\alpha) = \sum_{i,j} p_i^\alpha k_{ij}^{\alpha\beta} \quad \text{Flux}(B \rightarrow A|\beta) = \sum_{i,j} p_i^\beta k_{ij}^{\beta\alpha}. \quad (5)$$

4 Results

For AD, populations and MPFTs are measured in both directions between the two states as shown in Fig. 3. Here, we compare the results of the non-Markovian analysis with direct flux measurements to highlight the bias in the latter. Shown in Fig. 4 are the populations and passage times for the WE simulations. For reference, estimates from “brute force” molecular dynamics are also shown (black line) along with the 95% confidence estimates in the case of the MFPT (dashed black lines). The non-Markovian results are in excellent agreement with brute force estimates and well within range of statistical error, yet direct measurements show apparent bias. The reason for the large discrepancy is because direct measurement of p_α and p_β yields highly biased values (i.e. they differ substantially from brute force), despite the fact that state populations are reasonably well converged, thus biasing the MFPT estimates from Eq. 3.

In the methane system, the PMF and the MFPT from the unbound to bound state were measured as described above. Importantly, the passage times can be measured as function of the boundary definition for the unbound state from a single simulation. See the PMF inset of Fig. 5. Specifically, the bound state B was held fixed at a separation of 4 Å while the definition of the unbound state A was varied from 4 to 16 Å. The passage time from the bound to unbound state was measured in increments of 2 Å and is shown in Fig. 5. The result is in excellent agreement




Figure 3: A snapshot of the trajectory locations in (ϕ, ψ) space during a WE simulation of alanine-dipeptide. The state definitions are indicated by the A and B boxes. Red trajectories (α) were those last in state A, and blue trajectories (β) were those last in state B.

with [18]. Our variation of the boundary permits a physically reasonable selection near which the MFPT shows minimal sensitivity. The MFPT in the reverse direction follows a similar qualitative behavior (not shown).

For chignolin, two independent WE simulations were performed in addition to eight independent brute force simulations for comparison. The two WE simulations were initialized from different states: one from the native folded structure and one from an extended structure: see Fig. 6. Many of the well defined minima from the 2D PMF (Fig. 6) are in good agreement with those found in [23]. To illustrate the quality of sampling achieved by equilibrium WE, Fig. 7 below shows the probability distributions of the two chosen coordinates as a function of the simulation progression. In sharp contrast, brute force simulations remain trapped near the initial conformation given the same amount of total CPU time as the WE simulations ($1 \mu\text{s}$). The probability distributions obtained from 8 independent brute force simulation are shown in Fig. 8 and look similar to those in Fig. 7 for the (0%-1%) WE simulation range.

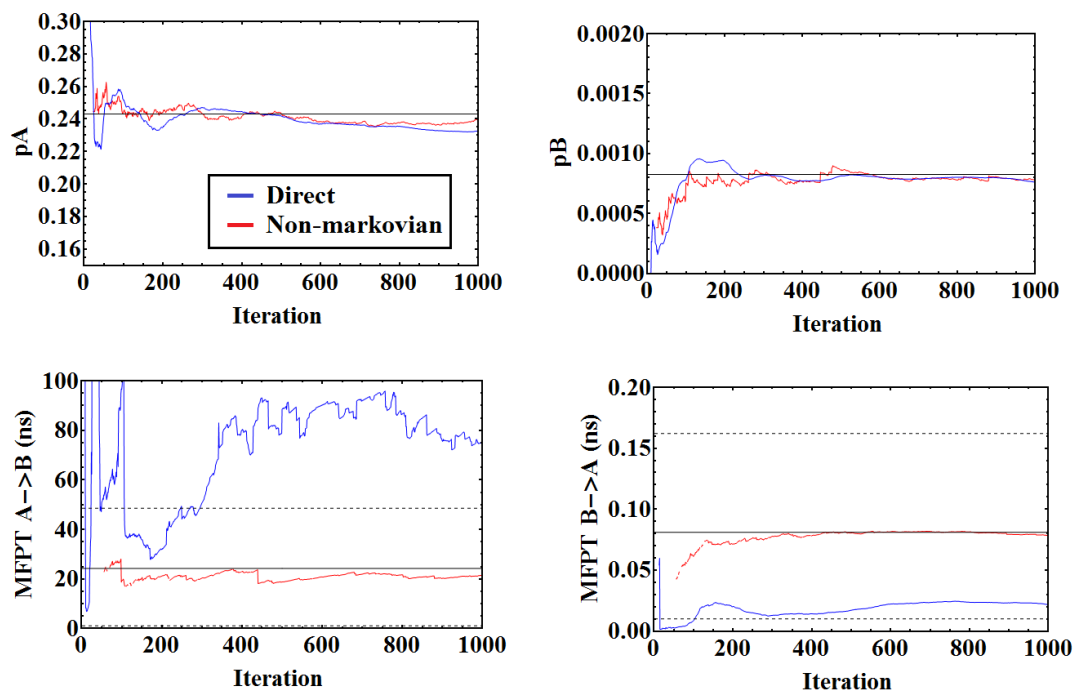


Figure 4: Populations and mean first passage times for alanine-dipeptide measured both directly and using the non-Markovian formalism. Solid lines indicate estimates from independent “brute force” simulations, with dashed lines indicating roughly a 95% confidence interval. The states are indicated in Fig. 3.




Figure 5: The mean first passage time for methane association (A to B) measured from WE simulation as a function of the boundary of state A. The inset displays the PMF along with the definitions of the unbound and bound states, indicated by A and B, respectively.

5 Concluding discussion

In summary, a generalized equilibrium weighted ensemble (WE) approach has been applied to measure equilibrium and kinetic properties in molecular systems. Most notably, it yields equilibrium and dynamical information in a single simulation. From a theoretical perspective, it is an exact approach and makes no Markovian assumption [20]. Because a true ensemble picture of equilibrium is simulated, arbitrary states can be chosen for analysis after the simulation is run. Not only does this provide great flexibility, but the WE trajectories themselves can guide the state definition process, a topic we hope to expand upon in future work. WE is also a parallel method characterized by essentially perfect scaling because communication among processors is intermittent. Although our current implementation employs user pre-defined bins, bins can be adaptively chosen [14, 20] as we will explore more fully in future work.

The formulation of a non-Markovian model of highly history-dependent bin populations is particularly important. Besides simplifying the calculation of non-Markovian kinetic quantities (i.e., rates between arbitrary states), the approach also removes a subtle but important bias that can occur when WE simulations are reweighted/resampled solely on the basis of equilibrium observables.




Figure 6: Two-dimensional PMF obtained from a WE simulation of chignolin as a function of the progress coordinates $X = r(\text{ASP3O-GLY7N})$ and $Y = r(\text{GLU5O-THR8N})$. The PMF was averaged over the second half of the simulation data. Simulations were initialized from the native and non-native (extended) states as indicated by the diamonds.




Figure 7: Evolution of the probability distributions $P(X)$ and $P(Y)$ for two WE simulations of chignolin. Blue curves are those initialized from the native state and red were initialized from the non-native. Shown are the distributions averaged over four periods of the WE simulation. At the end, both simulations approach the same distribution.

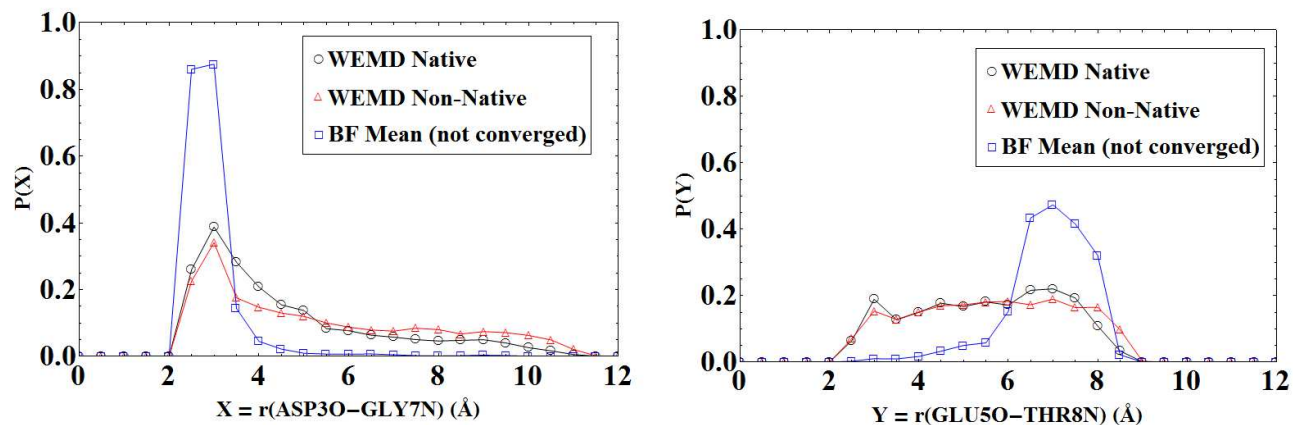


Figure 8: Comparison of WE to brute force (BF). The two WE simulations converge to one another despite starting in distant basins, whereas equal time BF simulations (8 trials) remained trapped near the initial conformation.

6 Acknowledgements

We thank Josh Adelman for insightful discussions, as well as the NSF for support through grants MCB-0643456 and MCB-1119091.

Appendix: Alternative solution for equilibrium bin populations based on minimal uncertainties

Although the steady-state solution of the rate matrix was used to calculate equilibrium bin populations in the results described above, other procedures are possible (because the equilibrium solution is over-determined given imperfect rate estimates from finite data). One alternative is to use the pairwise probability ratios from Eq. (1) to form clusters and ultimately to calculate all equilibrium populations, as described below. The basic strategy is to join bins (and fix their probability ratio) starting from pairs with lowest uncertainty in the corresponding rates — cf. Eq. (1). In general, separate clusters will be built up and then joined together.

One concrete procedure employed in our initial studies and in work to be presented elsewhere comprises the following series of steps:

1. Determine the pair of bins i and j whose estimated population ratio \hat{p}_i/\hat{p}_j has the minimum uncertainty among pairs not yet joined.
2. Join i and j in a cluster. If either or both are already in a cluster, join together whatever objects “own” bins i and j . Letting P_a and P_b denote the probabilities of the clusters owning bins i and j , the necessary ratio of cluster probabilities is $P_a/P_b = (P_a/\hat{p}_i)(\hat{p}_i/\hat{p}_j)(\hat{p}_j/P_b)$. This relation applies even if i was previously unclustered by setting $P_a = \hat{p}_i$ or analogously for j . Note that the fractional probability of a bin in a cluster is determined by the joining process.
3. Repeat from first step until all bins exhibiting any non-zero transition rate with another bin are joined into the smallest number of clusters possible. The total probability of a given cluster is not changed from the initial sum of probabilities of the constituent bins.

Once all bins are joined into a single cluster, the fractional bin populations for that cluster are the estimated equilibrium populations.

References

- [1] J. M. Berg, J. L. Tymoczko, and L. Stryer. *Biochemistry (5th ed.)*. Freeman, New York, 2002.
- [2] Daniel M Zuckerman. Equilibrium sampling in biomolecular simulations. *Annu Rev Biophys*, 40:41–62, Jun 2011. PMC Exempt - invited review.

- [3] R. H. Swendsen and J.-S. Wang. Replica Monte Carlo simulation of spin-glasses. *Phys. Rev. Lett.*, 57:2607–2609, 1986.
- [4] Lianqing Zheng, Mengen Chen, and Wei Yang. Random walk in orthogonal space to achieve efficient free-energy simulation of complex systems. *Proc Natl Acad Sci U S A*, 105(51):20227–20232, Dec 2008.
- [5] G. A. Huber and S. Kim. Weighted-ensemble Brownian dynamics simulations for protein association reactions. *Biophys. J.*, 70:97–110, 1996.
- [6] Peter G Bolhuis, David Chandler, Christoph Dellago, and Phillip L Geissler. Transition path sampling: throwing ropes over rough mountain passes, in the dark. *Annu Rev Phys Chem*, 53:291–318, 2002.
- [7] Rosalind J. Allen, Patrick B. Warren, and Pieter Rein ten Wolde. Sampling rare switching events in biochemical networks. *Physical Review Letters*, 94(1):018104, 2005.
- [8] Aryeh Warmflash, Prabhakar Bhimalapuram, and Aaron R. Dinner. Umbrella sampling for nonequilibrium processes. *The Journal of Chemical Physics*, 127(15):154112–8, 2007.
- [9] Eric Vanden-Eijnden and Maddalena Venturoli. Exact rate calculations by trajectory parallelization and tilting. *J Chem Phys*, 131(4):044120, Jul 2009.
- [10] David E Shaw, Paul Maragakis, Kresten Lindorff-Larsen, Stefano Piana, Ron O Dror, Michael P Eastwood, Joseph A Bank, John M Jumper, John K Salmon, Yibing Shan, and Willy Wriggers. Atomic-level characterization of the structural dynamics of proteins. *Science*, 330(6002):341–346, Oct 2010.
- [11] Nicolae-Viorel Buchete and Gerhard Hummer. Peptide folding kinetics from replica exchange molecular dynamics. *Phys Rev E Stat Nonlin Soft Matter Phys*, 77(3 Pt 1):030902, Mar 2008.
- [12] Xuhui Huang, Gregory R Bowman, Sergio Bacallado, and Vijay S Pande. Rapid equilibrium sampling initiated from nonequilibrium data. *Proc Natl Acad Sci U S A*, 106(47):19765–19769, Nov 2009.
- [13] Frank Noe, Christof Schtte, Eric Vanden-Eijnden, Lothar Reich, and Thomas R Weikl. Constructing the equilibrium ensemble of folding pathways from short off-equilibrium simulations. *Proc Natl Acad Sci U S A*, 106(45):19011–19016, Nov 2009.
- [14] Divesh Bhatt and Ivet Bahar. An adaptive weighted ensemble procedure for efficient computation of free energies and first passage rates. *J. Chem. Phys.*, 2012. In press.
- [15] Eric Darve and Ernest Ryu. *Innovations in Biomolecular Modeling and Simulations: Volume 1*, chapter Computing Reaction Rates in Bio-molecular Systems Using Discrete Macro-states, pages 138–206. Royal Society of Chemistry, 2012.

- [16] Bin W. Zhang, David Jasnow, and Daniel M. Zuckerman. Efficient and verified simulation of a path ensemble for conformational change in a united-residue model of calmodulin. *Proceedings of the National Academy of Sciences*, 104(46):18043–18048, 2007.
- [17] Divesh Bhatt and Daniel M. Zuckerman. Heterogeneous path ensembles for conformational transitions in semiatomistic models of adenylate kinase. *Journal of Chemical Theory and Computation*, 6(11):3527–3539, November 2010. PMCID: PMC3108504.
- [18] Matthew C. Zwier, Joseph W. Kaus, and Lillian T. Chong. Efficient explicit-solvent molecular dynamics simulations of molecular association kinetics: Methane/methane, Na⁺/Cl⁻, methane/benzene, and K⁺/18-Crown-6 Ether. *Journal of Chemical Theory and Computation*, 7(4):1189–1197, April 2011.
- [19] Divesh Bhatt, Bin W. Zhang, and Daniel M. Zuckerman. Steady state via weighted ensemble path sampling. *Journal of Chemical Physics*, 133:014110, 2010. PMCID: PMC2912933.
- [20] Bin W Zhang, David Jasnow, and Daniel M Zuckerman. The "weighted ensemble" path sampling method is statistically exact for a broad class of stochastic processes and binning procedures. *J Chem Phys*, 132(5):054107, Feb 2010. PMCID: PMC2830257.
- [21] Divesh Bhatt and Daniel M. Zuckerman. Beyond microscopic reversibility: Are observable nonequilibrium processes precisely reversible? *Journal of Chemical Theory and Computation*, 7(8):2520–2527, August 2011. PMCID in process.
- [22] Daniel M. Zuckerman. *Statistical Physics of Biomolecules: An Introduction*. CRC Press, Boca Raton, FL, 2010.
- [23] H. Okumura. Temperature and pressure denaturation of chignolin: Folding and unfolding simulation by multibaric-multithermal molecular dynamics method. *Proteins: Structure, Function, and Bioinformatics*, 80:2397–2416, 2012.